A First Look at GLAST LAT Absolute Timing

David A. Smith[1], J. Eric Grove[2], Denis Dumora[1], D. Patricia Sandora[2], and Eric J. Siskind[3]

[1] *Université Bordeaux 1 and CNRS/IN2P3, Centre d'Etudes Nucléaires de Bordeaux Gradignan, UMR 5797, Chemin du Solarium, BP120, 33175 Gradignan, France* `smith@cenbg.in2p3.fr`

[2] *Naval Research Laboratory, Code 7650, 4555 Overlook Ave. SW, Washington, DC 20375 USA*

[3] *NYCB Real-Time Computing, Inc., 18 Meudon Dr., Lattingtown, NY 11560 USA.*

28 February 2007, update 20 Mar 07.

Abstract

Feb 22 and 23 the GLAST LAT recorded cosmic ray muon tracks, with the Observatory GPS system providing absolute timing information. At the same time, the Bordeaux/NRL muon telescope was recording GPS times for a subset of the same muons. Correlation of the muon arrival times leads to strange and unexpected results. This write-up details the observations.

1 Update of previous results

In November 2006 we performed proof-of-principle tests, described in LAT-TD-08777-03. The VSC ("Virtual SpaceCraft") stood in for the real observatory, but without a GPS antenna, that is, the PPS provided to the LAT was *not* locked to absolute time. Figure 1 is an update of a concluding figure from that note. Recall that dT is the LAT time subtracted from the telescope GPS time. The most striking feature was a $\sim 2\,\mu$s/s drift, with a 24 ms offset, attributed to the undisciplined VSC GPS.

Two other curious features, at the few microsecond level, have since been discovered to be due to an analysis error (the elapsed time used for the correction was an integer number of seconds). Figure 1, bottom right, was, in consequence, a near-flat distribution $\pm 1\,\mu$s wide. Since correcting the error, it is better described by a normal distribution, with an RMS width of 160 ns. Further, the bottom middle plot seemed sinusoidal. It is in fact piecewise linear. The slope of the 4 segments shown are -2.46, $+3.01$, -2.46, and -8.0 ns/s, respectively. The turn-arounds at 570 and 1500 s correspond to excursion of $\pm 1.3\,\mu$s from the mean value of dT. It may be due to the "Bordeaux" GPS PPS precision being updated as satellite positions vary.

Further small improvements to the analysis code were made, concerning the detector geometry, and the handling of the deviation of the LAT system clock from the nominal 20 MHz. The improvements due to these latter improvements are minor.

Figure 2 carries the analysis a step forward. The scaler rate `dTicks`, described in the next section, is constant throughout the half-hour run (bottom left), at a value of $80 \times 50 = 4\,\mu$s/s deviation from 20 MHz.

In the following, we will see that some disturbing "features" exist in the new data acquired since LAT-Observatory integration.

2 Geometry, Code, and Data

As shown in figures 3 and 4, the scintillator paddle positions are different than they were in November. The November muon rate of 4.5 Hz is now 2.4 Hz, since the telescope now looks at a larger zenith angle.

Table 1 lists the 8 data runs: 4 with config A and 4 with config B. For each set, 2 with GPS lock and 2 with GPS not locked. The run numbers are 14191,2,3,4 and 14215,6,7,8. LAT muon data rates are 480 Hz and each half-hour run contains about $880,000$ events. The standalone VME GPS runs contain about 4500 events each, taken concurrently with the LAT runs. LAT data are compressed in B0-9-0.

Table 1 lists the data runs taken. The "difference from 20 MHz",`dTicks`, is $20,000,000$ - `TicksPerSecond`, where `TicksPerSecond` is the difference between the values of `TimeToneCurrentGemTimeTicks` and `TimeTonePreviousGemTimeTicks`, which are the scaler readings at PPS arrival times, for the

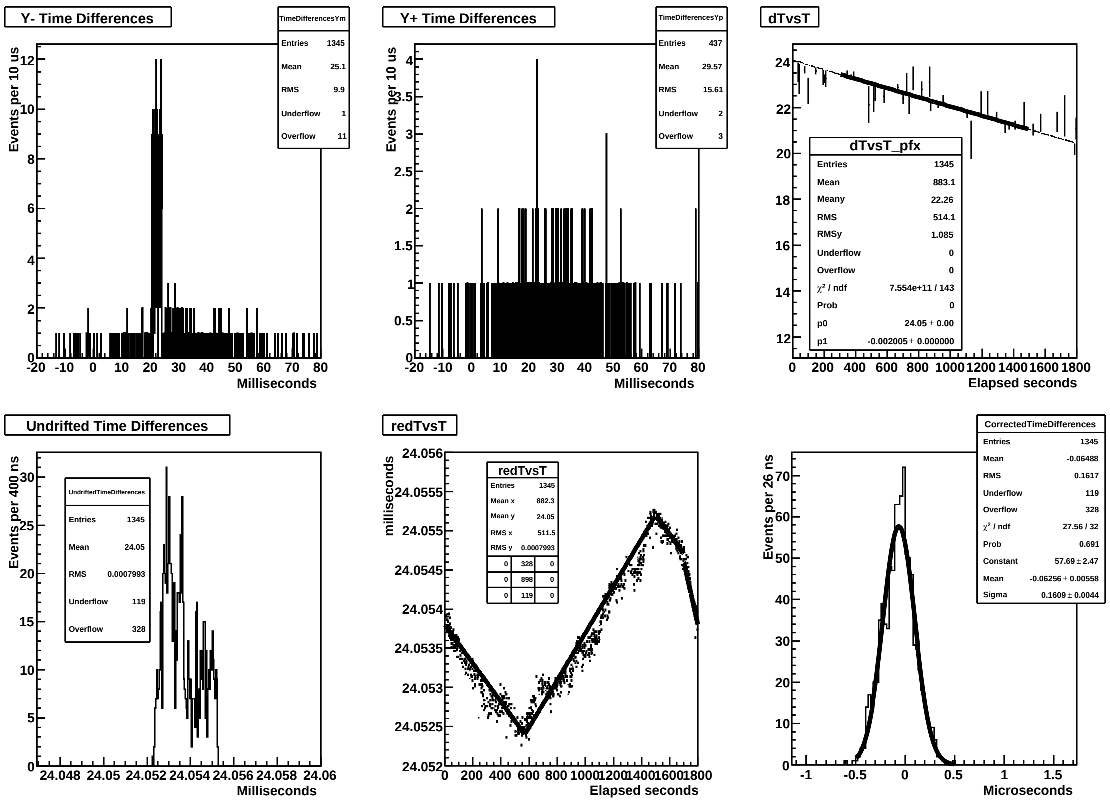

Figure 1: Run 77013003 from November 2006, re-visited, with code improvements. Top left: Time differences for LAT muon tracks striking both muon scintillators. The spike at $dT \simeq 24$ ms is the signal, the other events are accidental coincidences . Top middle: same, but for a "virtual" telescope on the opposite side of the LAT. Only accidentals are observed. Top right: dT versus the elapsed time – a steady drift of $2.009\,\mu$s per second is observed. Bottom left: after subtracting the drift. Bottom middle: Again dT versus elapsed time, but after drift subtraction. A smaller piece-wise linear drift appears, at -2.4 ns/s reaching $-1.3\,\mu$s below the mean, then switching to $+3.0$ ns/s reaching $1.3\,\mu$s above the mean. It is possibly the GPS adjusting the PPS rate as the satellite fixes vary. Bottom right: After subtracting the piecewise linear function the root-mean-square variance of the distribution is 160 ns.

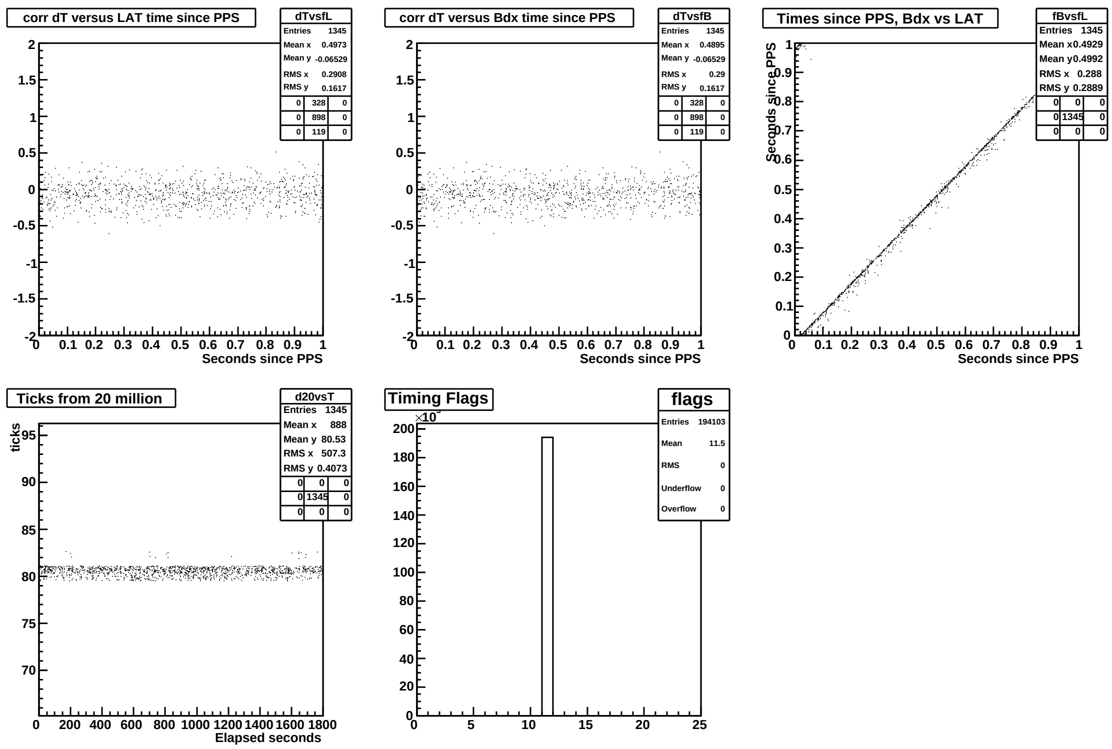

Figure 2: Run 13003 from November, continued. Top Left: Corrected dT, i.e. the same events as in the last frame of figure 1, plotted versus the fraction of a second between two PPS arrivals, for the LAT. Top Middle and Right: The fraction obtained from the telescope GPS is strongly correlated with the LAT fraction. Bottom Left: The quantity $(20,000,000-\mathrm{PPSticks})$, where PPSticks is the difference in the number of "20 MHz" scaler ticks between the Current event and the Previous event, for the muon coincidences, plotted versus run elapsed time. Bottom middle: The various Context flags concatenated into a single word. For this run, all events have the value 11, meaning that ContextLsfTimeTimeToneMissingGps, Previous and Current, are always set to 1 for all events, while all other flags are 0.

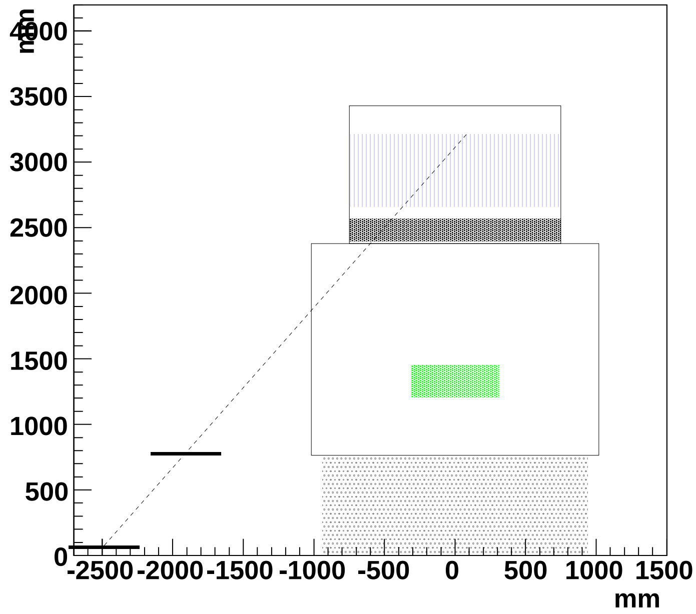

Figure 3: New layout. The LAT mounted on the spacecraft is higher above the floor. The muon telescope is a little farther away. The top scintillator paddle is offset more than for the November runs.

current and the preceding events. If one were to assume that the nominal 20 MHz oscillator driving the scaler really gave twenty million ticks per second, then one would induce an error on the fractional part of the second, the magnitude of which would be roughly `TicksPerSecond` times 50 ns.

Here follows the algorithm used to convert the LAT variables obtained from the root SvacTuple into MET times. The "Context" (and all other) variables in the SvacTuple are described in Anders' documentation,

ftp://ftp-glast.slac.stanford.edu/glast.u33/u01/svac_workshop/Svac_ntuple_description/new/html/inde

The quite long variable names have been abbreviated. For example, `ContextLsfTimeTimeToneCurrentGemTimeTicks` becomes `TTCGemTimeTicks` and, similarly, `ContextLsfTimeTimeTonePrevious` becomes `TTP`. The algorithm is then

```
   // Rollover offset for the 25 bits GEM counter:
double RollOver = 33554432.0;

   TicksPerSecond[iLAT] = TTCGemTimeTicks[iLAT] - TTPGemTimeTicks[iLAT];
if (TTCGemTimeTicks[iLAT] < TTPGemTimeTicks[iLAT])
TicksPerSecond[iLAT] += RollOver;

   EvtTicksSincePPS[iLAT] = TimeTicks[iLAT] - TTCGemTimeTicks[iLAT];
if (TimeTicks[iLAT] < TTCGemTimeTicks[iLAT])
EvtTicksSincePPS[iLAT] += RollOver;

   fraction[iLAT] =
((double)EvtTicksSincePPS[iLAT]) / ((double) TicksPerSecond[iLAT]) ;
```

Then, to calculate dT, the time difference between the LAT and standalone GPS times, we use

```
Double_t dT =
(double)(SecsBdx[iBdx]-TTCTSecs[iLAT]) + FracBdx[iBdx]-fraction[iLAT] ;
```

where `Bdx` is short for Bordeaux, the home of the standalone GPS. Time coincidences were one full second off: we have added 1 second to Bordeaux times for the February data. A GD

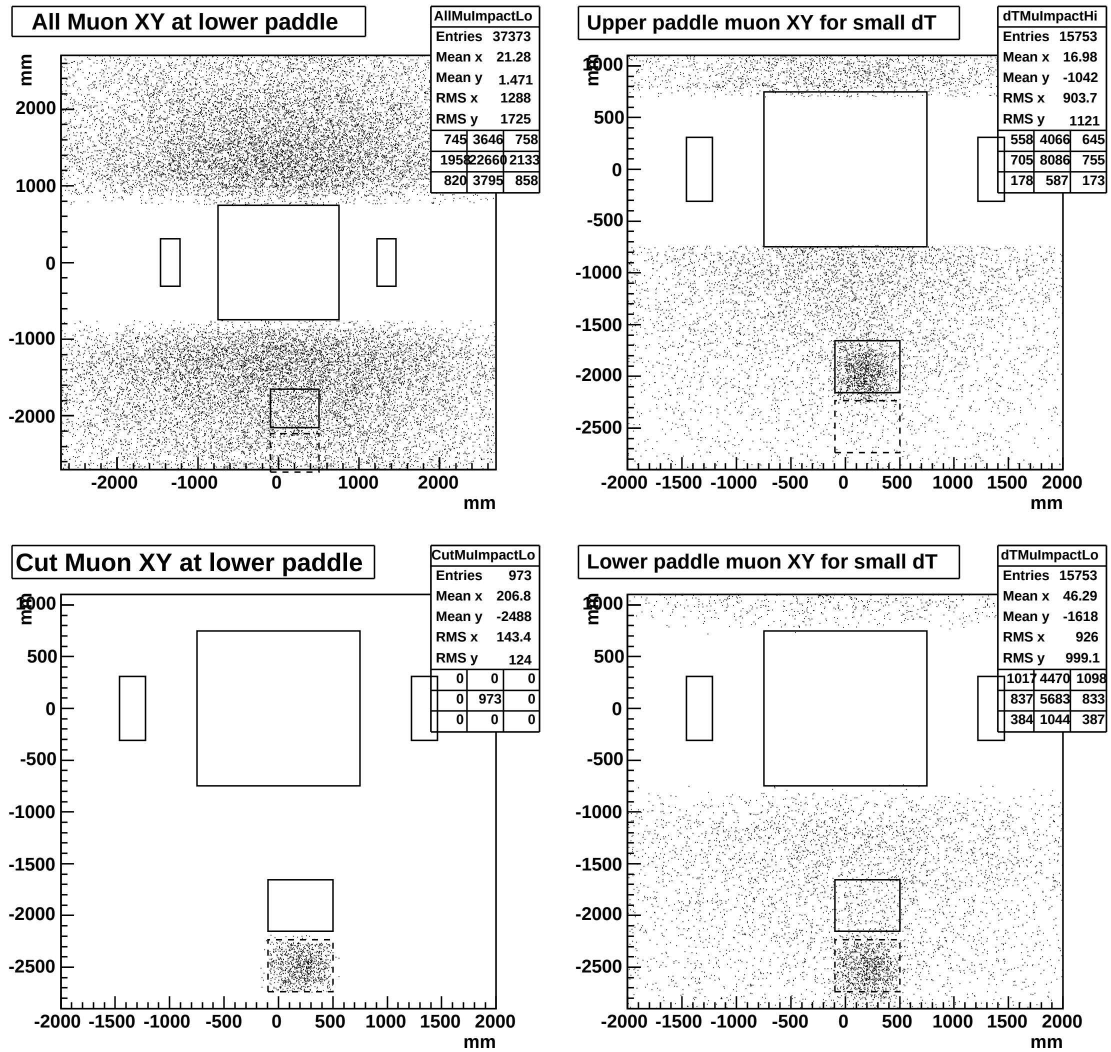

Figure 4: New layout, continued. Bird's eye view of extrapolated muon impact points using reconstructed LAT tracks, for run 77014215. The large square is the LAT. The small solid (dashed) square is the upper (lower) scintillator paddle. The rectangles at left and at right are GBM BGO counters. Top left: In the plane of the lower paddle (i.e., almost on the HiBay floor). To keep file sizes small, only events reaching the $\pm$Y sides are kept. Here, only 1 in 10 events is shown, for clarity. Top right: Requiring a $|dT| < 10$ ms time difference between the LAT and the standalone GPS clock, in the upper paddle plane. A clear event excess appears. Bottom right: Same, for lower paddle. Bottom left: Impact points in the plane of the lower panel when the track is required to pass through the upper panel (no timing cut).

LAT run	GPS lock?	Config/Side	Diff. from 20 MHz ($\sim$ 50 ns ticks)	Observed dT drift	Wraparound?
77014191	Yes	2/B	110	$-3.4\,\mu$s/s	Yes
77014192	Yes	2/B	110	$-3.4\,\mu$s/s	Yes
77014193	No	2/B	306	$-13\,\mu$s/s	No
77014194	No	2/B	306	$-13\,\mu$s/s ?	No
77014215	Yes	1/A		$-3.4\,\mu$s/s	Yes
77014216	Yes	1/A	115 to 125	$-3.4\,\mu$s/s	Yes
77014217	No	1/A	-70 to -60	$+7\,\mu$s/s	No
77014218	No	1/A	-60 to -50	$+7\,\mu$s/s	No

Table 1: The 8 data sets acquired February 22 and 23, 2007.

test engineer told us that an S/C software initialization that would have corrected this was not performed.

3 Results

Figures 5 through 12 pretty much speak for themselves.

When the S/C GPS lock is on, dT decreases from 0 to -1 ms over about 290 seconds ($-3.4\,\mu$s/s), then wraps back to zero.

When the lock is off, the drift goes to tens of milliseconds, without wrapping around. The sign of the drift is negative for runs '193 and '194 (config 2, side B) and is positive for runs '217 and '218. The sign of the drift correlates with the sign of the difference from twenty million. The magnitude of the drift does too.

Figure 6, bottom middle, shows a "tic tac toe" board with the plot occupancy: $5 + 3$ events have `dTicks` values outside the y-range of the scatter plot.

Figures 13 and 14 reveal that the outliers cluster at values of $\pm 10,000$ and $-20,000$ ticks, that is, ± 0.5 ms and -1 ms. Examination of the muon coincidence events shows that these outlier occur at the -1 ms "wrap-around" times of the sawtooth in the GPS lock runs, lasting for 2 or 3 seconds, as the timing circuit settles in. The 10k and 20k outliers do not occur in the runs without GPS lock.

In the 10 days following the initial release of this memo, we (mainly EJS) have analysed the available documentation of the spacecraft UDL circuits (Up- and DownLink, or "oodle"), and of the LAT GEM. A preliminary interpretation of the observations is that there could be a PPS polarity inversion between two subsystems. The UDL-to-GEM PPS signal has a 250 ms/750 ms duty cycle that does not manifest itself, eliminating that explanation. On the other hand, the GPS-to-UDL PPS signal polarity may be inverted. Verification of this hypothesis, and options for mitigation, are currently underway.

3.0.1 Why it's unlikely to be the muon telescope or the data analysis

N. Johnson pointed out that $2^{32} \times 50$ ns $= 215$ seconds, which is in the same ballpark as the 290 s period of the $0 > dT > -1$ ms sawtooth observed for the runs with GPS lock. He then asked whether either the Bordeaux VME GPS code, or the data analysis code might not have some 32-bit integer being loaded with overly large values.

DAS re-verified his LAT analysis code with this in mind, and thinks not. The largest numbers are the MET seconds, and these are clearly uncorrupted. The GEM scaler ticks 'only' go up to 2^{25}. For the fractional part of the second, `double` precision is used.

To check against this sort of error in a big piece of the LAT ground analysis software, we fed the coincidence event list for run 14215 to a routine written by Anders Borgland, `checkDigi.C`. Recall that the "digi root" file is a re-formatting of the "lsf" files that arrive at SLAC, while the "svacTuple" root file is a derivative of the "recon" file, where tracks and energy have been reconstructed. These are large codes, where accidents could in principle happen.

Figure 15 clearly reconstructs the sawtooth pattern, so the problem does *not* arrive in the digi-to-svacTuple chain.

Figure 16 breaks the sawtooth down by PPS intervals, similar to figure 2. Nothing striking appears.

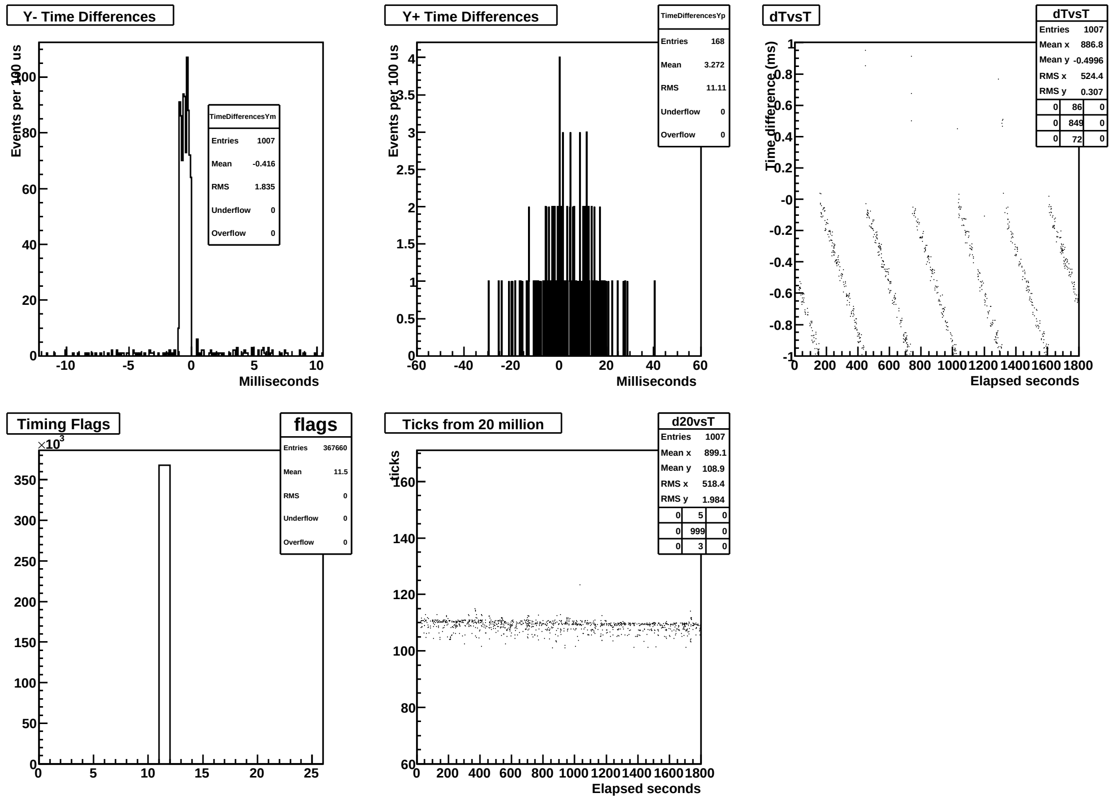

Figure 5: Run 77014191. Top row, as in figure 1. Bottom row, as in figure 2. Isn't the 1 ms wrap-around weird?

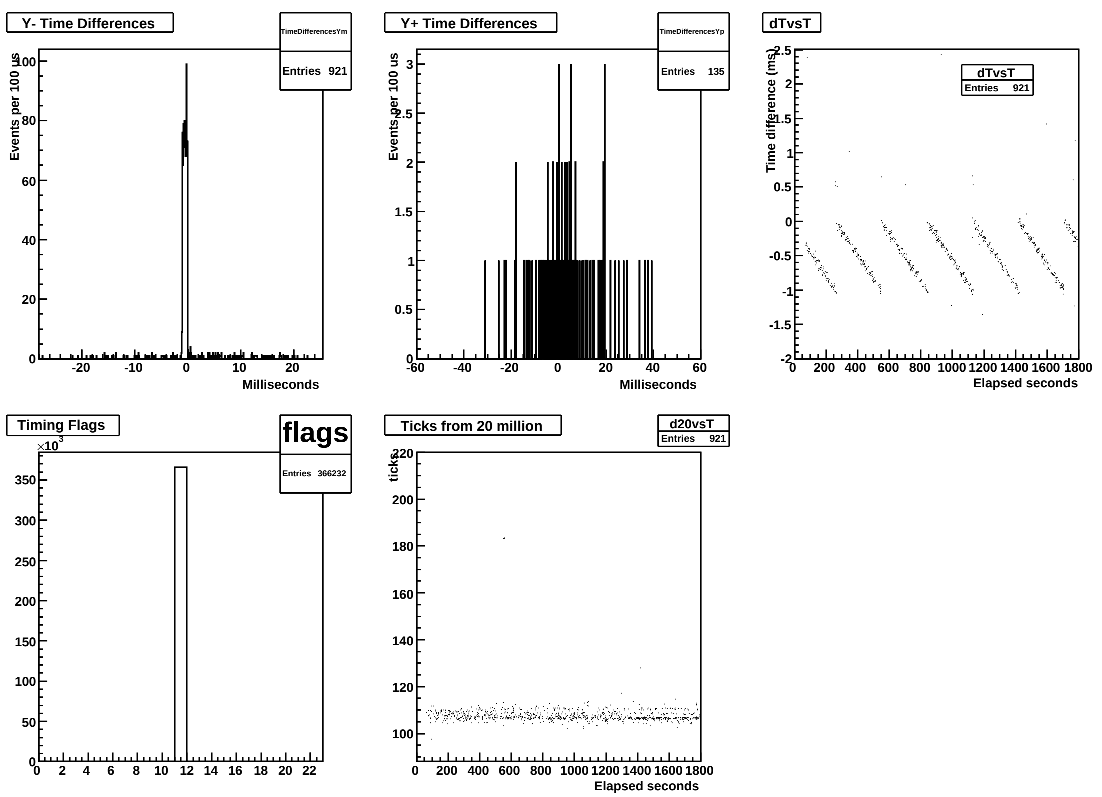

Figure 6: Run 77014192. Like preceding figure. Slope is 1 ms in about 290 seconds, which is 3.4 μs per second. Note that 110 ticks away from 20 MHz is like a 5.5 μs drift, remarkably close to the observed drift.

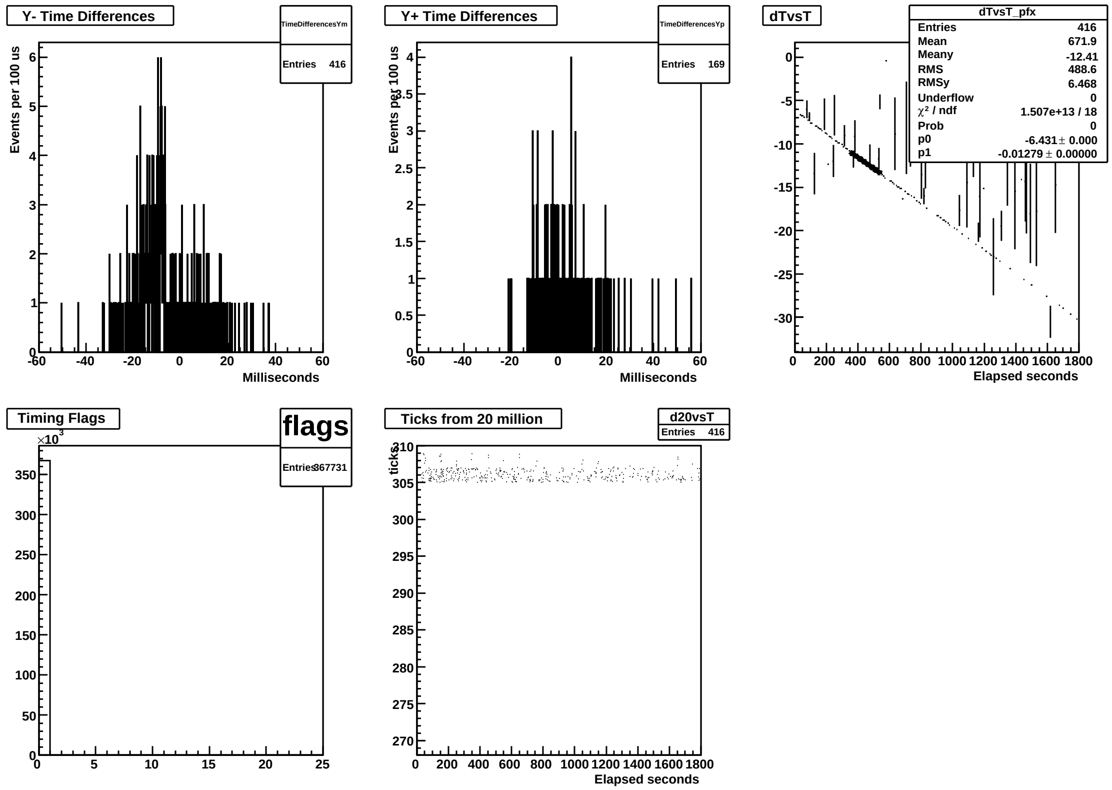

Figure 7: Run 77014193. Like preceding figure. With GPS unlocked, wraparound no longer happens. The top right plot (dT vs elapsed time) looks ratty due to the Profile binning used to perform fit to get the drift rate. Next plot has the prettier non-profile as well.

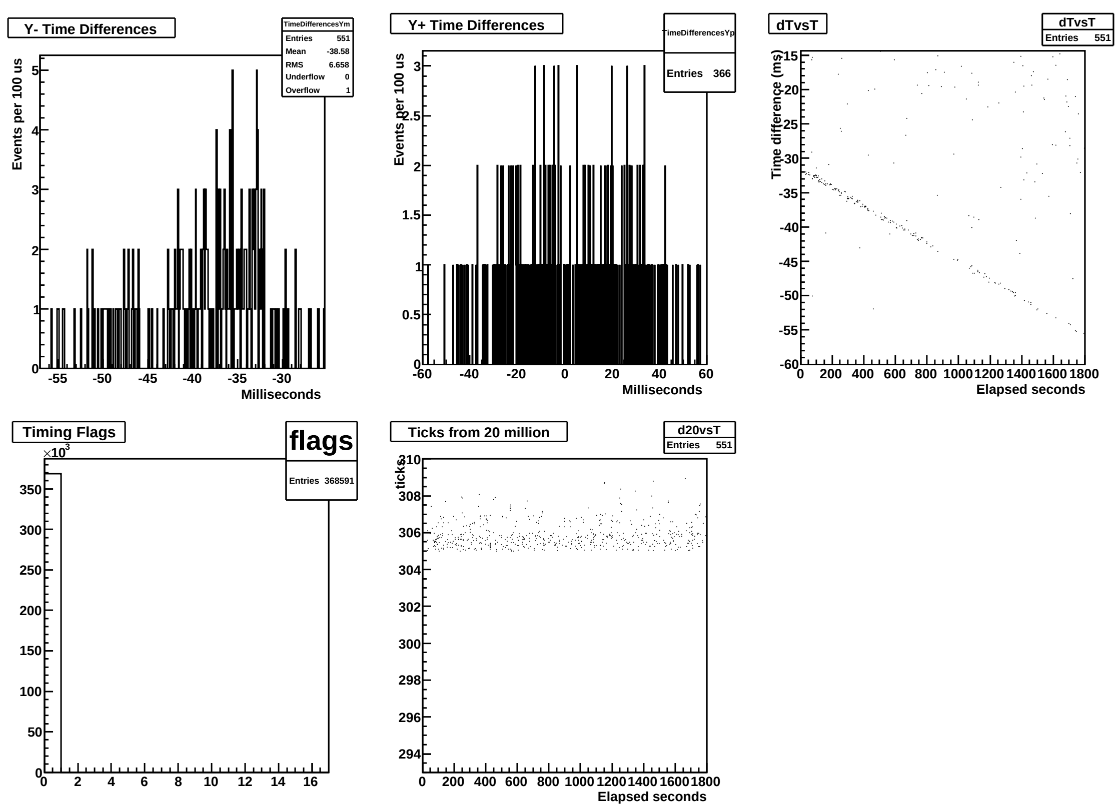

Figure 8: Run 77014194. Like preceding figure. The slope from -29 ms in the preceding run just keeps going, all the way out to -55 ms.

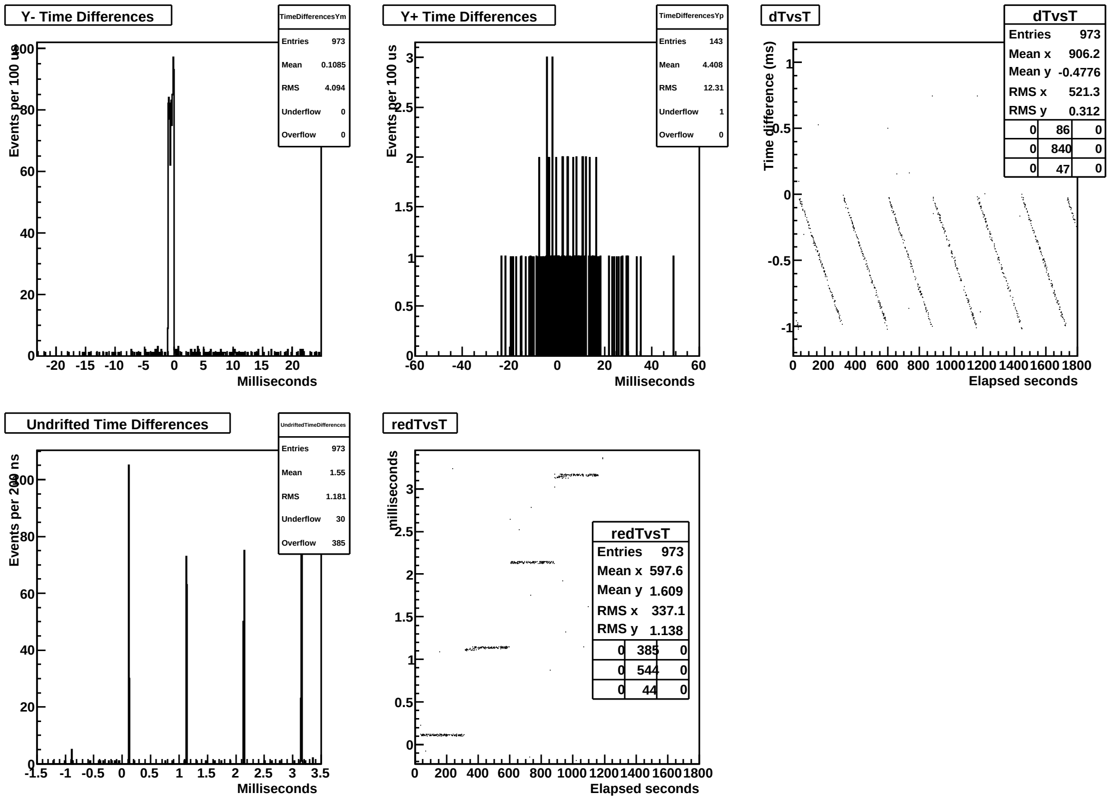

Figure 9: Run 77014215, same frames as for figure 5. The 1 ms wrap-around is back! The period is about 290 seconds.

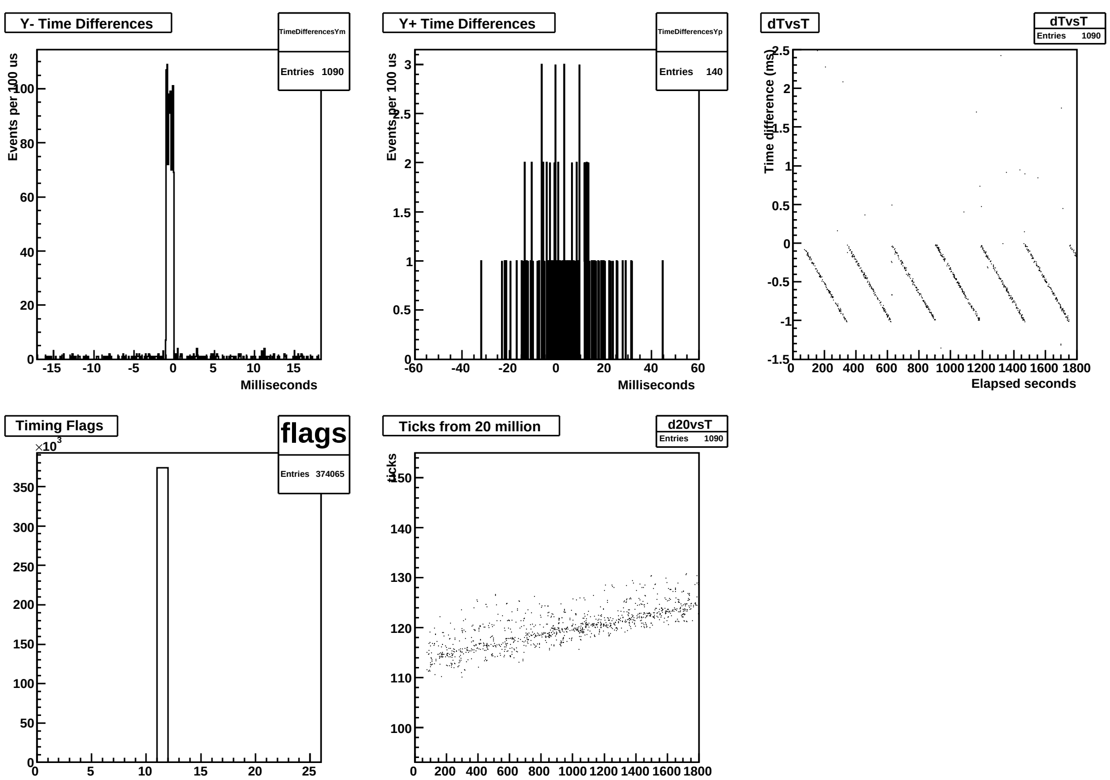

Figure 10: Run 77014216, same as preceding figure.

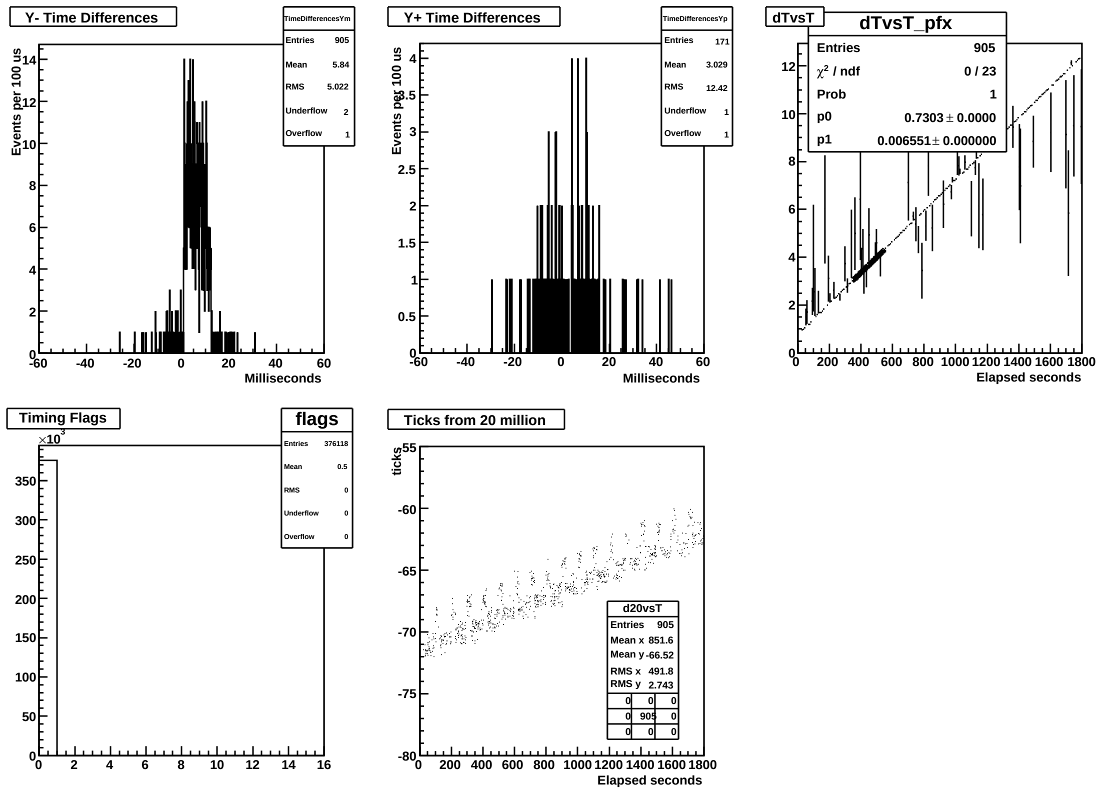

Figure 11: Run 77014217, more of the same.

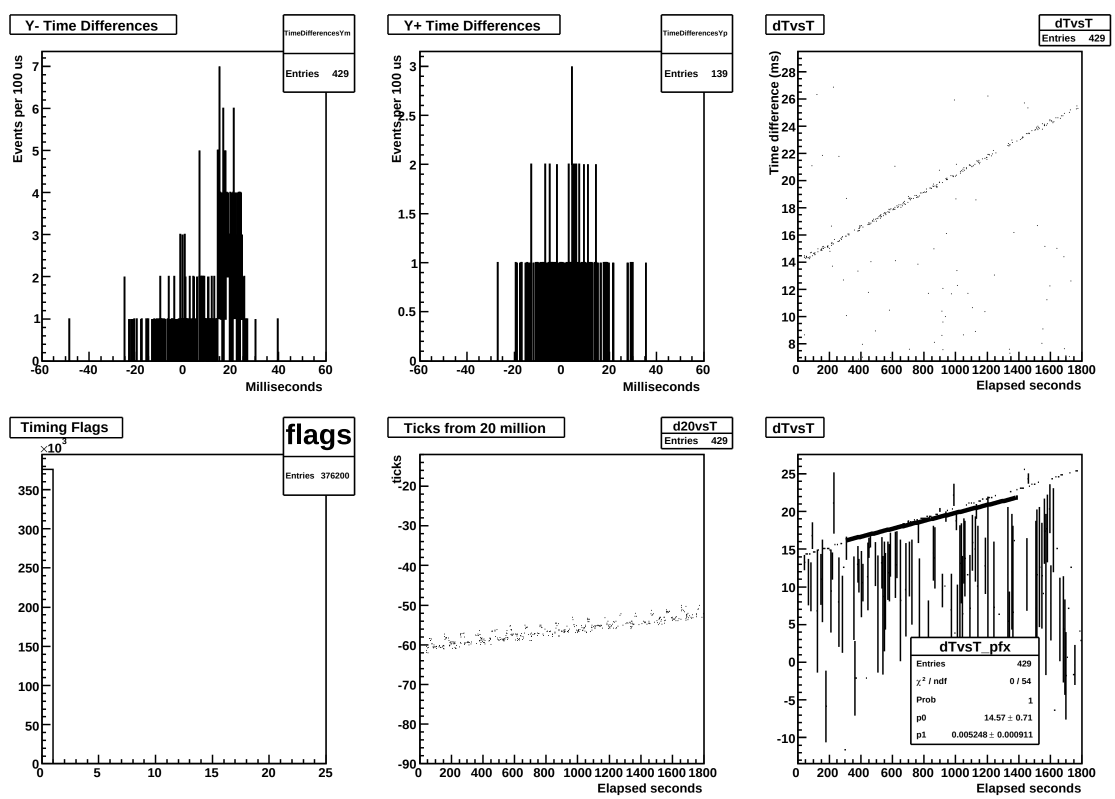

Figure 12: Run 77014218, same as it ever was.

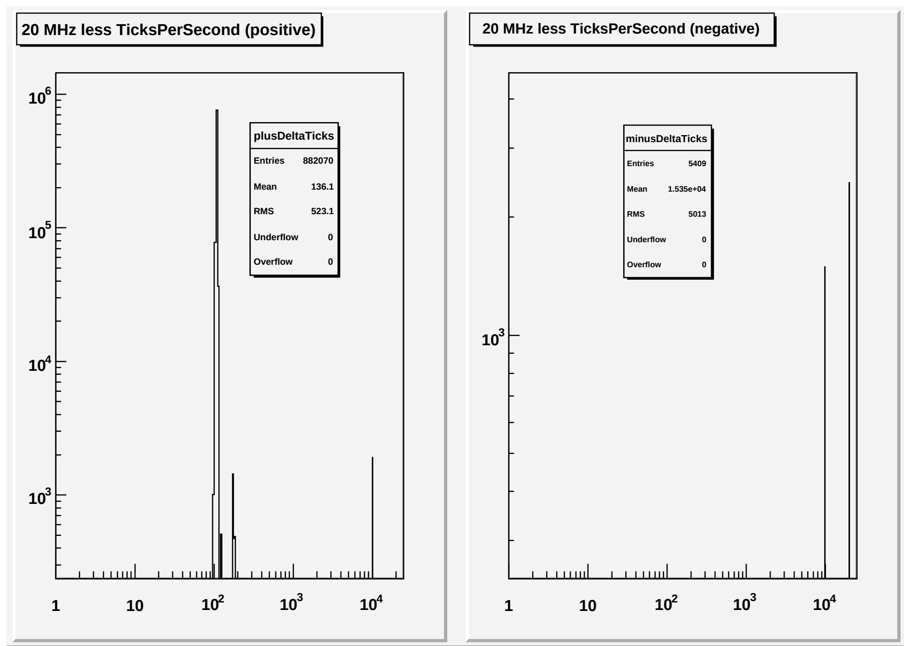

Figure 13: Run 77014215. For all events (not just muon coincidences), `dTicks`, which is the deviation of the number of (nominally) 20 MHz scaler ticks between two successive PPS signals from 20, 000, 000. Left: Positive values of $dTicks$. Right: Negative values

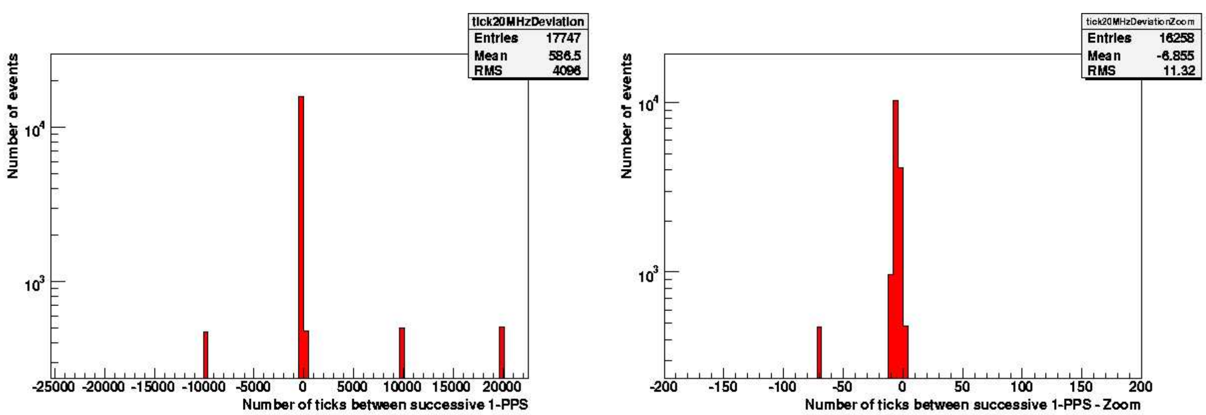

Figure 14: Run 77014210 (same S/C and LAT configuration as 14214, that is, GPS locked. `dTicks`, as in previous figure and text, plotted differently. Left: all events. Right: Zoom on central region. *Credit: A. Borgland*

Could it be in the Bordeaux VME GPS acquisition code? Here are some arguments against that:

1. The acquisition code is very simple. The integer number of seconds is clearly uncorrupted, since LAT and Bordeaux agree to within 10's of milliseconds. The fractional part is a `double` built from a sequence of short integers.

2. The basic code has not changed since the Crab optical pulsar was detected with this hardware+software, as described in the ApJ, reference [14] of LAT-TD-08777.

3. Our symptom changes significantly for the 4 LAT+S/C configurations, while the Bordeaux configuration is constant.

4. Figure 3 of LAT-TD-08777 shows sub-microsecond agreement between two Bordeaux-type GPS's.

DAS fears that the problem lies elsewhere...

3.0.2 Further Checks of "Bordeaux" Timestamps

We recently acquired another GPS, the idea being that if we were somehow mis-using the Symmetricom devices, that the mistake would appear when comparing with completely different hardware. The new GPS is an *RF Solutions* model LS-40EB, mounted on their LS-40EVALR1 evaluation kit (168 euros, from Farnell). It was the cheapest, easy-to-use device we could find having a PPS output. Plugging it into the serial port of a PC permits visualisation of satellites being used and of the expected precision ("PDOP"), but the GPS automatically locks and provides PPS signals without the laptop.

We used the *RF Solutions* PPS output to trigger the VME "time capture", expecting to see VME GPS times very close to an integer number of seconds. Figure 17 shows the results for an overnight run of 50 ks (13.9 hours). For 96% of the events, the times are within $\pm 0.5\,\mu$s, better than *RF Solutions* claim of $\pm 1\,\mu$s PPS accuracy. The Symmetricom dates are, on average, 580 ns early compared to the PPS arrivals, roughly consistent with the ~ 100 meter cable run of the Trimble Accutime antenna on the roof at the far end of the building. (The new antenna was only a few meters away, outside of the room with the VME crate.)

1979 events have time deviations reaching 3.5 ms. These occur 12k seconds into the run, and again between 36k and 39k seconds. Presumably, the *RF Solutions* antenna stopped seeing enough satellites, given that the building blocks the northern sky. The Trimble antenna on the roof sees more. Furthermore, the Symmetricom circuit is designed to remain accurate during GPS unlock, whereas the LS-40EB documentation says that the PPS output is disabled for 3 minutes after loss of satellite fix. Apparently it keeps on ticking, but wrongly.

The 9 underflow events, and 11 overflow deviations beyond the $< +3.5$ ms just mentioned, are different: They have values of -325, ~ -100, and $\sim +100$ ms in the first few thousand seconds of the run, do not occur closely spaced in time, and are not obviously associated with satellite lock loss. A minor mystery.

3.0.3 A glance at the future

Moving away from GPS etc hardware... the Science Tools used to make pulsar light curves get the gamma ray times from a fits file, which is a copy of a subset of variables stored in the MeritTuple root file. Anders Borgland wrote the code that calculates the time stamp from the GEM variables, using an algorithm very similar to the one used in this note.

Figure 18 shows that he did it just fine. It's the same sawtooth as before, but this time, using the LAT times from the standard Science Tools fits file, created from the standard MeritTuple.

Stuff to not forget DAS says that the events were not time-ordered in the files, at least before the update of the decompression code, and Anders disagrees. Follow through...

Presumably, spacecraft telemetry includes the typical generic GPS information (PDOP, satellite almanach, and so forth). It would be nice if the ISOC control room had access to this information to allow us to scrutinize timestamp quality should we ever see pulsar phase anomalies.

Acknowledgements: Anders Borgland and Gregg Thayer are providing tons of help.

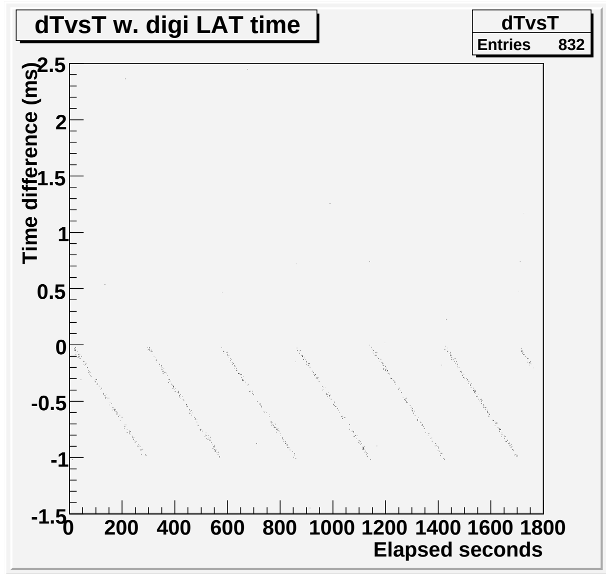

Figure 15: Different ground analysis code using the digi.root file for Run 77014215. Same as it ever was.

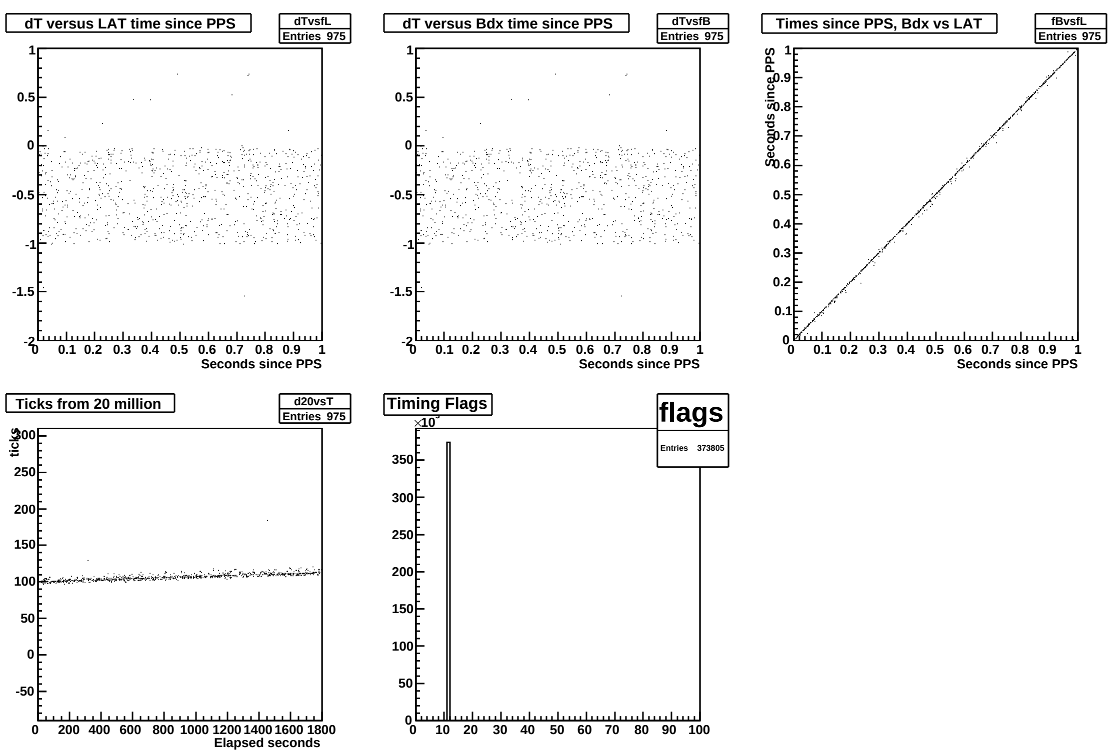

Figure 16: (Back to the svacTuple code) For run 77014215, the -1 ms sawtooth folded with the fraction of a second between the current and the upcoming PPS.

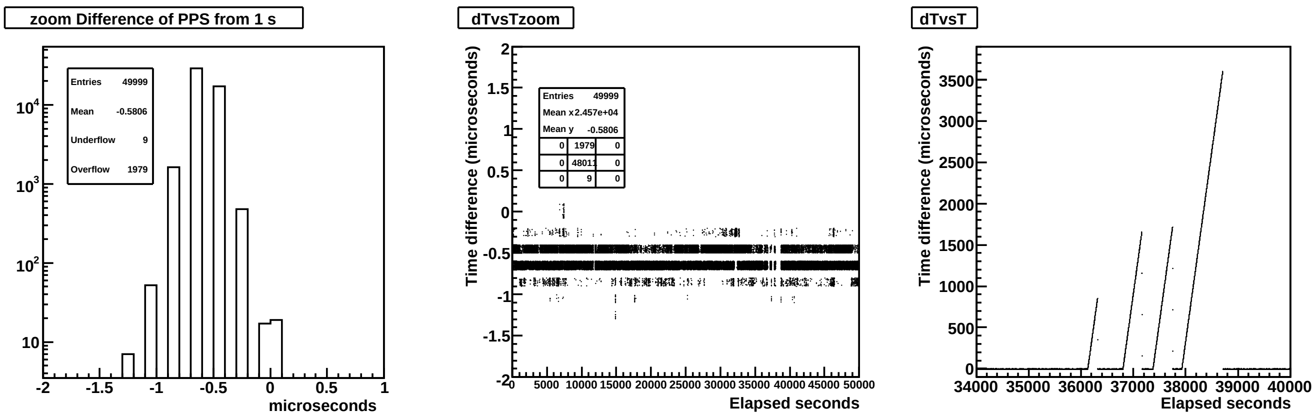

Figure 17: The deviation of the Symmetricom VME GPS times from an integer number of seconds, when the VME "time capture" is triggered by the PPS output of the *RF Solutions* GPS. Left: Histogram of values. The mean of -580 ns is roughly consistent with the cable run of the Trimble antenna used for the Symmetricom. The dispersion is better than the $\pm 1\,\mu$s claimed by *RF Solutions*. Middle: Values versus elapsed time. It appears that satellite lock was lost briefly after 12k seconds, and again between 36k and 39k seconds. Right: Zoom on the long GPS-unlock period.

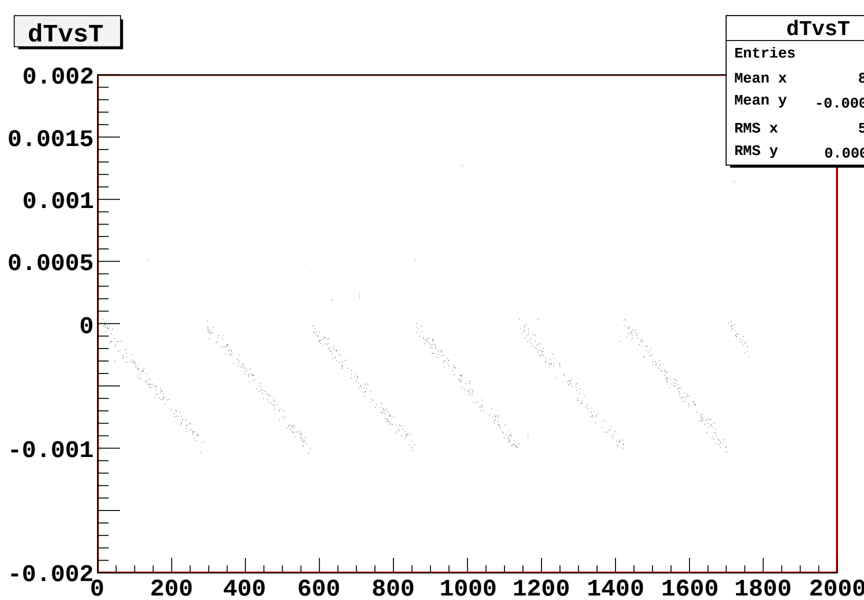

Figure 18: Same sawtooth from run 14215, but made with the dates from the Science Tools FITS file, derived from the MeritTuple root file. Anders Borgland calculates the MeritTuple times, and Thierry Reposeur made this plot.